

# Strong Consistency of the Good-Turing Estimator

Aaron B. Wagner

Coordinated Science Laboratory

Univ. of Illinois at Urbana-Champaign  
and School of ECE, Cornell University

wagner@ece.cornell.edu

Pramod Viswanath

Coordinated Science Laboratory

and ECE Dept.

Univ. of Illinois at Urbana-Champaign

pramodv@uiuc.edu

Sanjeev R. Kulkarni

EE Dept.

Princeton University  
kulkarni@princeton.edu

**Abstract**—We consider the problem of estimating the total probability of all symbols that appear with a given frequency in a string of i.i.d. random variables with unknown distribution. We focus on the regime in which the block length is large yet no symbol appears frequently in the string. This is accomplished by allowing the distribution to change with the block length. Under a natural convergence assumption on the sequence of underlying distributions, we show that the total probabilities converge to a deterministic limit, which we characterize. We then show that the Good-Turing total probability estimator is strongly consistent.

## I. INTRODUCTION

The problem of estimating the underlying probability distribution from an observed data sequence arises in a variety of fields such as compression, adaptive control, and linguistics. The most familiar technique is to use the empirical distribution of the data, also known as the type. This approach has a number of virtues. It is the maximum likelihood (ML) distribution, and if each symbol appears frequently in the string, then the law of large numbers guarantees that the estimate will be close to the true distribution.

In some situations, however, not all symbols will appear frequently in the observed data. One example is a digital image with the pixels themselves, rather than bits, viewed as the symbols [1]. Here the size of the alphabet can meet or exceed the total number of observed symbols, i.e., the number of pixels in the image. Another example is English text. Even in large corpora, many words will appear once or twice or not at all [2]. This makes estimating the distribution of English words using the type ineffective. This problem is particularly pronounced when one attempts to estimate the distribution of bigrams, or pairs of words, since the number of bigrams is evidently the square of the number of words.

To see that the empirical distribution is lacking as an estimator for the probabilities of uncommon symbols, consider the extreme situation in which the alphabet is infinite and we observe a length- $n$  sequence containing  $n$  distinct symbols [3]. The ML estimator will assign probability  $1/n$  to the  $n$  symbols that appear in the string and zero probability to the rest. But common sense suggests that the  $(n+1)$ st symbol in the sequence is very likely to be one that has not yet appeared. It seems that the ML estimator is overfitting the data. Modifications to the ML estimator such as the Laplace “add one” and the Krichevsky-Trofimov “add half” [4] have been proposed as remedies, but these only alleviate the problem [3].

In collaboration with Turing, Good [5] proposed an estimator for the probabilities of rare symbols that differs considerably from the ML estimator. The Good-Turing estimator has been shown to work well in practice [6], and it is now used in several application areas [3]. Early theoretical work on the estimator focused on its bias [5], [7], [8]. Recent work has been directed toward developing confidence intervals for the estimates using central limit theorems [9], [10] or concentration inequalities [11], [12]. Orlitsky, Santhanam, and Zhang [3] showed that the estimator has a pattern redundancy that is small but not optimal. None of these works, however, have shown that the estimator is strongly consistent.

We show that the Good-Turing estimator is strongly consistent under a natural formulation of the problem. We consider the problem of estimating the total probability of all symbols that appear  $k$  times in the observed string for each nonnegative integer  $k$ . For  $k = 0$ , this is the total probability of the unseen symbols, a quantity that has received particular attention [7], [13]. Estimating the total probability of all symbols with the same empirical frequency is a natural approach when the symbols appear infrequently so that there is insufficient data to accurately estimate the probabilities of the individual symbols. Although the total probabilities are themselves random, we show that under our model they converge to a deterministic limit, which we characterize. Note that if the alphabet is small and the block length is large, then the problem effectively reduces to the usual probability estimation problem since it is unlikely that multiple symbols will have the same empirical frequency.

It is known that the Good-Turing estimator performs poorly for high-probability symbols [3], but this is not a problem since the ML estimator can be employed to estimate the probabilities of symbols that appear frequently in the observed string. We therefore focus on the situation in which the symbols are unlikely, meaning that they have probability  $O(1/n)$ . We allow the underlying distributions to vary with the block length  $n$  in order to maintain this condition, and we assume that, properly scaled, these distributions converge. This model is discussed in detail in the next section, where we also describe the Good-Turing estimator. In Section III, we establish the convergence of the total probabilities. Section IV uses this convergence result to show strong consistency of the Good-Turing estimator. Some comments regarding how to estimate other quantities of interest are made in the final section.

## II. PRELIMINARIES

Let  $(\Omega_n, \mathcal{F}_n, P_n)$  be a sequence of probability spaces. We do not assume that  $\Omega_n$  is finite or even countable. Our observed string is a sequence of  $n$  symbols drawn i.i.d. from  $\Omega_n$  according to  $P_n$ . Note that the alphabet and the underlying distribution are permitted to vary with  $n$ . This allows us to model the situation in which the block length is large while the number of occurrences of some symbols is small.

### A. Total Probabilities

For each nonnegative integer  $k$ , let  $A_k^n$  denote the set of symbols in  $\Omega_n$  that appear exactly  $k$  times in the string of length  $n$ . We call

$$\xi_k^n := P_n(A_k^n)$$

the *total probability* of symbols that appear  $k$  times.

Of course, for  $k \geq 1$ ,  $\xi_k^n$  is simply the sum of the probabilities of the symbols with frequency  $k$ . On the other hand,  $A_0^n$  will be uncountable if  $\Omega_n$  is.

We view  $\xi_k^n$  as a random probability distribution on the nonnegative integers. Our goal is to estimate this distribution.

### B. The Good-Turing Estimator

The Good-Turing estimator is normally viewed as an estimator for the probabilities of the individual symbols. Let  $\varphi_k^n = |A_k^n|$  denote the number of symbols that appear exactly  $k$  times in the observed sequence. The basic Good-Turing estimator assigns probability

$$\frac{(k+1)\varphi_{k+1}^n}{n\varphi_k^n}$$

to each symbol that appears  $k \leq n-1$  times [5]. The case  $k=n$  must be handled separately, but this case is unimportant to us since under our model it is unlikely that only one symbol will appear in the string.

This formula can be naturally viewed as a total probability estimator since the  $\varphi_k^n$  in the denominator is merely dividing the total probability equally among the  $\varphi_k^n$  symbols that appear  $k$  times. Thus the Good-Turing total probability estimator assigns probability

$$\zeta_k^n := \frac{(k+1)\varphi_{k+1}^n}{n}$$

to the aggregate of symbols that have appeared  $k$  times for each  $k$  in  $\{0, \dots, n-1\}$ . As a convention, we shall always assign zero probability to the set of symbols that appear  $n$  times

$$\zeta_n^n := 0.$$

Like  $\xi_k^n$ ,  $\zeta_k^n$  is a random probability distribution on the nonnegative integers.

As a total probability estimator,  $\zeta_k^n$  is not ideal. For one thing,  $\zeta_k^n$  can be positive even when  $A_k^n$  is empty, in which case  $\xi_k^n$  is clearly zero. A similar problem arises when estimating the probabilities of individual symbols, and modifications to the basic Good-Turing estimator have been proposed to avoid it [5]. But we shall show that even the basic form of the Good-Turing estimator is strongly consistent for total probability estimation.

### C. Shadows

The distributions of the total probability,  $\xi_k^n$ , and the Good-Turing estimator,  $\zeta_k^n$ , are unaffected if one relabels the symbols in  $\Omega_n$ . This fact makes it convenient in what follows to consider the probabilities assigned by  $P_n$  without reference to the labeling of the symbols.

*Definition 1:* Let  $X_n$  be a random variable on  $\Omega_n$  with distribution  $P_n$ . The *shadow* of  $P_n$  is defined to be the distribution of the random variable  $P_n(\{X_n\})$ .

As an example, if  $\Omega_n = \{a, b, c\}$  and

$$P_n(\{a\}) = P_n(\{b\}) = \frac{1}{2} P_n(\{c\}) = \frac{1}{4},$$

then the shadow of  $P_n$  would be uniform over  $\{1/4, 1/2\}$ . If  $P_n$  is itself uniform, then its shadow is deterministic. Note that the discrete entropy of a distribution only depends on the distribution through its shadow. We will write  $P_n(X_n)$  as a shorthand for  $P_n(\{X_n\})$  in what follows.

For finite alphabets, specifying the shadow is equivalent to specifying the unordered components of  $P_n$ , viewed as a probability vector. This is clearly seen in the above example, since the shadow is uniformly distributed over  $\{1/4, 1/2\}$  if and only if the underlying distribution has two symbols with probability 1/4 and one with probability 1/2.

If  $P_n$  has a continuous component, then the shadow will have a point mass at zero equal to the probability of this component. The shadow reveals nothing more about the continuous component than its total probability, but we shall have no need for such information. Indeed, the distributions of both  $\xi_k^n$  and  $\zeta_k^n$  depend on  $P_n$  only through its shadow.

### D. Unlikely Symbols

To prove strong consistency, we assume that the scaled profiles,  $n \cdot P_n(X_n)$ , converge to a nonnegative random variable  $Y$  with distribution  $Q$ . This implies, in particular, that asymptotically almost every symbol has probability  $O(1/n)$  and therefore appears  $O(1)$  times in the sequence on average. As an example, if  $P_n$  is a uniform distribution over an alphabet of size  $n$ , then the scaled shadow,  $n \cdot P_n(X_n)$ , equals one a.s. for each  $n$  (and hence it converges in distribution). More complicated examples can be constructed by quantizing a fixed density more and more finely to generate the sequence of distributions.

## III. TOTAL PROBABILITY CONVERGENCE

Before considering the performance of the Good-Turing estimator, we study the asymptotics of the total probabilities themselves. Under our assumption that the scaled shadows converge, we show that the total probabilities converge almost surely to a deterministic Poisson mixture.

*Proposition 1:* The random distribution  $\xi^n$  converges to

$$\lambda_k := \int_0^\infty \frac{y^k \exp(-y)}{k!} dQ(y) \quad k = 0, 1, 2, \dots$$

in  $L^1$  almost surely as  $n \rightarrow \infty$ .

We prove this result by first showing that the mean of  $\xi^n$  converges to  $\lambda$  and then proving concentration around the

mean. To show convergence of the mean, it is convenient to make several definitions. Let

$$g_k^n(y) = \binom{n}{k} \left(\frac{y}{n}\right)^k \left(1 - \frac{y}{n}\right)^{n-k}$$

and

$$g_k(y) = \frac{y^k \exp(-y)}{k!}.$$

Since

$$\binom{n}{k} \frac{1}{n^k} \rightarrow \frac{1}{k!} \quad \text{as } n \rightarrow \infty$$

and

$$\left(1 + \frac{y_n}{n}\right)^n \rightarrow \exp(y) \quad \text{if } y_n \rightarrow y,$$

it follows that for all sequences  $y_n \rightarrow y$ ,  $g_k^n(y_n) \rightarrow g_k(y)$ . Note also that  $g_k^n(y) \leq 1$  if  $0 \leq y \leq n$  by the binomial theorem. Let

$$C^n = \{\omega \in \Omega_n : P_n(\omega) > 0\}$$

and note that  $C^n$  is countable for each  $n$ .

*Lemma 1:* For all nonnegative integers  $k$ ,

$$\lim_{n \rightarrow \infty} E[\xi_k^n] = \lambda_k.$$

*Proof:* We shall show that

$$E[\xi_k^n] = E[g_k^n(nP_n(X_n))]. \quad (1)$$

First consider the case  $k \geq 1$ . Here

$$\begin{aligned} \xi_k^n &= P_n(A_k^n \cap C^n) \\ &= \sum_{\omega \in C^n} 1(\omega \in A_k^n) P_n(\omega) \end{aligned}$$

so by monotone convergence

$$\begin{aligned} E[\xi_k^n] &= \sum_{\omega \in C^n} \binom{n}{k} P_n(\omega)^k (1 - P_n(\omega))^{n-k} P_n(\omega) \\ &= \sum_{\omega \in C^n} g_k^n(nP_n(\omega)) P_n(\omega) \\ &= E[g_k^n(nP_n(X_n)) 1(X_n \in C^n)] \\ &= E[g_k^n(nP_n(X_n))]. \end{aligned}$$

Next consider the case  $k = 0$ . Here

$$\begin{aligned} \xi_0^n &= P_n(A_0^n) \\ &= P_n(A_0^n \cap C^n) + P_n(A_0^n - C^n) \\ &= \sum_{\omega \in C^n} 1(\omega \in A_0^n) P_n(\omega) + P_n(\Omega_n - C^n). \end{aligned}$$

So again by monotone convergence,

$$\begin{aligned} E[\xi_0^n] &= \sum_{\omega \in C^n} (1 - P_n(\omega))^n P_n(\omega) + P_n(\Omega_n - C^n) \\ &= \sum_{\omega \in C^n} g_0^n(nP_n(\omega)) P_n(\omega) + P_n(\Omega_n - C^n) \\ &= E[g_0^n(nP_n(X_n)) 1(X_n \in C^n)] \\ &\quad + E[g_0^n(nP_n(X_n)) 1(X_n \notin C^n)] \\ &= E[g_0^n(nP_n(X_n))]. \end{aligned}$$

This establishes (1). Since  $nP_n(X_n)$  converges in distribution to  $Y$ , we can create a sequence of random variables  $\{Y_n\}_{n=1}^\infty$  such that  $Y_n$  has the same distribution as  $nP_n(X_n)$  and  $Y_n$  converges to  $Y$  almost surely [14, Theorem 4.30]. Then

$$g_k^n(Y_n) \rightarrow g_k(Y) \quad \text{a.s.}$$

Since  $g_k^n(Y_n) \leq 1$  a.s., the bounded convergence theorem implies

$$\begin{aligned} \lim_{n \rightarrow \infty} E[g_k^n(Y_n)] &= E[g_k(Y)] \\ &= \int_0^\infty g_k(y) dQ(y) = \lambda_k. \end{aligned}$$

□

*Lemma 2:* For all nonnegative integers  $k$ ,

$$\lim_{n \rightarrow \infty} |\xi_k^n - E[\xi_k^n]| = 0 \quad \text{a.s.}$$

*Proof:* Let

$$B^n = \left\{ \omega \in \Omega_n : P_n(\omega) \geq \frac{1}{n^{3/4}} \right\}$$

and note that  $|B^n| \leq n^{3/4}$ . Then let

$$\tilde{\xi}_k^n = P_n(A_k^n \cap B^n),$$

and note that

$$|\xi_k^n - E[\xi_k^n]| \leq |(\xi_k^n - \tilde{\xi}_k^n) - E[\xi_k^n - \tilde{\xi}_k^n]| + \tilde{\xi}_k^n + E[\tilde{\xi}_k^n].$$

Now if we change one symbol in the underlying sequence, then  $\xi_k^n - \tilde{\xi}_k^n$  can change by at most  $2/n^{3/4}$ . By the Azuma-Hoeffding-Bennett concentration inequality [15, Corollary 2.4.14], it follows that for all  $\epsilon > 0$

$$\Pr \left( |(\xi_k^n - \tilde{\xi}_k^n) - E[\xi_k^n - \tilde{\xi}_k^n]| \geq \epsilon \right) \leq 2 \exp \left[ -\frac{\epsilon^2 \sqrt{n}}{8} \right].$$

Since the right-hand side is summable over  $n$ , this implies that

$$|(\xi_k^n - \tilde{\xi}_k^n) - E[\xi_k^n - \tilde{\xi}_k^n]| \rightarrow 0 \quad \text{a.s.}$$

Now

$$\tilde{\xi}_k^n = \sum_{\omega \in B^n} P_n(\omega) 1(\omega \in A_k^n)$$

so

$$\begin{aligned} E[\tilde{\xi}_k^n] &= \sum_{\omega \in B^n} P_n(\omega) \binom{n}{k} (P_n(\omega))^k (1 - P_n(\omega))^{n-k} \\ &\leq \sum_{\omega \in B^n} \binom{n}{k} (P_n(\omega))^k (1 - P_n(\omega))^{n-k}. \end{aligned}$$

But

$$\begin{aligned} \binom{n}{k} (P_n(\omega))^k (1 - P_n(\omega))^{n-k} \\ = \exp \left[ -n \left( H \left( \frac{k}{n} \right) + D \left( \frac{k}{n} \middle\| P_n(\omega) \right) \right) \right], \end{aligned}$$

where  $H(\cdot)$  denotes the binary entropy function and  $D(\cdot \parallel \cdot)$  denotes binary Kullback-Leibler divergence, both with natural

logarithms [16, Theorem 12.1.2]. For all sufficiently large  $n$ ,  $k/n < 1/n^{3/4}$ , which implies that for all  $\omega \in B^n$ ,

$$D\left(\frac{k}{n}\middle\|P_n(\omega)\right) \geq D\left(\frac{k}{n}\middle\|\frac{1}{n^{3/4}}\right).$$

This gives

$$\begin{aligned} \binom{n}{k}(P_n(\omega))^k(1 - P_n(\omega))^{n-k} \\ \leq \binom{n}{k} \left(\frac{1}{n^{3/4}}\right)^k \left(1 - \frac{1}{n^{3/4}}\right)^{n-k}, \end{aligned}$$

so

$$E[\tilde{\xi}_k^n] \leq n^{3/4} \binom{n}{k} \left(\frac{1}{n^{3/4}}\right)^k \left(1 - \frac{1}{n^{3/4}}\right)^{n-k}.$$

Since

$$\binom{n}{k} \leq \frac{n^k}{k!},$$

this implies

$$E[\tilde{\xi}_k^n] \leq \frac{n^{(k+3)/4}}{k!} \left(1 - \frac{1}{n^{3/4}}\right)^{n-k}. \quad (2)$$

Now the right-hand side tends to zero as  $n \rightarrow \infty$ , so

$$\lim_{n \rightarrow 0} E[\tilde{\xi}_k^n] = 0.$$

In fact, the right-hand side of (2) is summable over  $n$ . By Markov's inequality,

$$\Pr(\tilde{\xi}_k^n > \epsilon) \leq \frac{E[\tilde{\xi}_k^n]}{\epsilon},$$

this implies that  $\tilde{\xi}_k^n \rightarrow 0$  a.s. The conclusion follows.  $\square$

*Proof of Proposition 1:* It follows from Lemmas 1 and 2 that for each  $k$ ,

$$\lim_{n \rightarrow \infty} \xi_k^n = \lambda_k \quad \text{a.s.}$$

That is,  $\xi^n$  converges pointwise to  $\lambda$  with probability one. The strengthening to  $L^1$  convergence follows from Scheffé's theorem [17, Theorem 16.12], but we shall give a self-contained proof since it is brief. Observe that with probability one,

$$\begin{aligned} 0 &= \sum_{k=0}^{\infty} [\lambda_k - \xi_k^n] \\ &= \sum_{k=0}^{\infty} [\lambda_k - \xi_k^n]^+ - \sum_{k=0}^{\infty} [\lambda_k - \xi_k^n]^-, \end{aligned}$$

where  $[\cdot]^+$  and  $[\cdot]^-$  represent the positive and negative parts, respectively. Thus

$$\sum_{k=0}^{\infty} |\lambda_k - \xi_k^n| = 2 \sum_{k=0}^{\infty} [\lambda_k - \xi_k^n]^+ \quad \text{a.s.}$$

But  $[\lambda_k - \xi_k^n]^+$  converges pointwise to 0 a.s. and is less than or equal to  $\lambda_k$ . The dominated convergence theorem then implies that

$$\lim_{n \rightarrow \infty} \sum_{k=0}^{\infty} [\lambda_k - \xi_k^n]^+ = 0 \quad \text{a.s.}$$

#### IV. STRONG CONSISTENCY

The key to showing strong consistency is to establish a convergence result for the Good-Turing estimator that is analogous to Proposition 1 for the total probabilities.

*Proposition 2:* The random distribution  $\zeta^n$  converges to  $\lambda$  in  $L^1$  almost surely as  $n \rightarrow \infty$ .

The desired strong consistency follows from this result and Proposition 1.

*Theorem 1:* The Good-Turing total probability estimator is strongly consistent, i.e.,

$$\lim_{n \rightarrow \infty} \sum_{k=0}^n |\xi_k^n - \zeta_k^n| = 0 \quad \text{a.s.}$$

*Proof:* We have

$$\sum_{k=0}^n |\xi_k^n - \zeta_k^n| \leq \sum_{k=0}^{\infty} |\xi_k^n - \lambda_k| + \sum_{k=0}^{\infty} |\lambda_k - \zeta_k^n|.$$

We now let  $n \rightarrow \infty$  and invoke Propositions 1 and 2.  $\square$

The proof of Proposition 2 parallels that of Proposition 1 in the previous section. In particular, we first show that the mean of  $\zeta^n$  converges to  $\lambda$  and then establish concentration around the mean.

*Lemma 3:* For all nonnegative integers  $k$ ,

$$\lim_{n \rightarrow \infty} E[\zeta_k^n] = \lambda_k.$$

*Proof:* We shall show that

$$E[\zeta_k^n] = E[g_k^{n-1}((n-1)P_n(X_n))]. \quad (3)$$

First consider the case  $k \geq 1$ . Here

$$\zeta_k^n = \sum_{\omega \in C^n} \frac{k+1}{n} 1(\omega \in A_{k+1}^n).$$

So by monotone convergence,

$$\begin{aligned} E[\zeta_k^n] &= \sum_{\omega \in C^n} \frac{k+1}{n} \binom{n}{k+1} (P_n(\omega))^{k+1} (1 - P_n(\omega))^{n-k-1} \\ &= \sum_{\omega \in C^n} \binom{n-1}{k} (P_n(\omega))^k (1 - P_n(\omega))^{n-k-1} P_n(\omega) \\ &= \sum_{\omega \in C^n} g_k^{n-1}((n-1)P_n(\omega)) P_n(\omega) \\ &= E[g_k^{n-1}((n-1)P_n(X_n)) 1(X_n \in C^n)] \\ &= E[g_k^{n-1}((n-1)P_n(X_n))]. \end{aligned}$$

Next consider the case  $k = 0$ . Here

$$\begin{aligned} \zeta_0^n &= \frac{1}{n} |A_1^n| \\ &= \frac{1}{n} |A_1^n \cap C^n| + \frac{1}{n} |A_1^n - C^n| \\ &= \frac{1}{n} \sum_{\omega \in C^n} 1(\omega \in A_1^n) + \frac{1}{n} |A_1^n - C^n|. \end{aligned}$$

$\square$

Again invoking monotone convergence,

$$\begin{aligned}
E[\zeta_0^n] &= \frac{1}{n} \sum_{\omega \in C^n} \binom{n}{1} P_n(\omega) (1 - P_n(\omega))^{n-1} \\
&\quad + P_n(\Omega_n - C^n) \\
&= \sum_{\omega \in C^n} g_0^{n-1}((n-1)P_n(\omega)) P_n(\omega) \\
&\quad + P_n(\Omega_n - C^n) \\
&= E[g_0^{n-1}((n-1)P_n(X_n)) \mathbf{1}(X_n \in C^n)] \\
&\quad + E[g_0^{n-1}((n-1)P_n(X_n)) \mathbf{1}(X_n \notin C^n)] \\
&= E[g_0^{n-1}((n-1)P_n(X_n))].
\end{aligned}$$

This establishes (3). Following the reasoning in the proof of Lemma 1, this implies

$$\lim_{n \rightarrow \infty} E[\zeta_k^n] = E[g_k(Y)] = \lambda_k$$

for all  $k$ .  $\square$

*Lemma 4:* For all nonnegative integers  $k$ ,

$$\lim_{n \rightarrow \infty} |\zeta_k^n - E[\zeta_k^n]| = 0 \quad \text{a.s.}$$

*Proof:* Observe that if we alter one symbol in the underlying i.i.d. sequence, then  $\zeta_k^n$  will change by at most  $2(k+1)/n$ . As in the proof of Lemma 2, the Azuma-Hoeffding-Bennett concentration inequality [15, Corollary 2.4.14] then implies that

$$\Pr(|\zeta_k^n - E[\zeta_k^n]| > \epsilon) \leq 2 \exp \left[ -\frac{\epsilon^2 n}{8(k+1)^2} \right].$$

Since the right-hand side is summable over  $n$ , the conclusion follows.  $\square$

*Proof of Proposition 2:* The result follows from Lemma 3, Lemma 4, and Scheffé's theorem [17, Theorem 16.12] as in the proof of Proposition 1.  $\square$

## V. SHADOW ESTIMATION

Proposition 1 shows that the total probabilities converge to a deterministic limit, which is a function of the limit of the scaled shadows,  $Q$ . In fact, the total probabilities converge to a Poisson mixture, with  $Q$  being the mixing distribution. The functional form of the Poisson distribution enables us to create a simple function of the observed string, the Good-Turing estimator, that has the same limit as the total probabilities. In particular, we can consistently estimate the total probabilities

without having to explicitly estimate  $Q$ .

In general, such a shortcut might not be available. It is of interest therefore to study how to estimate  $Q$  itself from the observed string. With an estimator for  $Q$ , one could create a "plug-in" estimator for other quantities of interest.

## ACKNOWLEDGMENT

This research was supported in part by the Army Research Office under grant DAAD19-00-1-0466 and the National Science Foundation under grants CAREER-0237549, ITR-0325924, and CCR-0312413.

## REFERENCES

- [1] A. Orlitsky, N. P. Santhanam, and J. Zhang, "Universal compression of memoryless sources over unknown alphabets," *IEEE Trans. Inf. Theory*, vol. 50, no. 7, pp. 1469–81, July 2004.
- [2] B. Efron and R. Thisted, "Estimating the number of unseen species: How many words did Shakespeare know?" *Biometrika*, vol. 63, no. 3, pp. 435–47, 1976.
- [3] A. Orlitsky, N. P. Santhanam, and J. Zhang, "Always Good Turing: Asymptotically optimal probability estimation," *Science*, vol. 302, pp. 427–31, Oct. 2003.
- [4] R. Krichevsky and V. Trofimov, "The performance of universal encoding," *IEEE Trans. Inf. Theory*, vol. 27, no. 2, pp. 199–207, Mar. 1981.
- [5] I. J. Good, "The population frequencies of species and the estimation of population parameters," *Biometrika*, vol. 40, no. 3/4, pp. 237–64, 1953.
- [6] W. Gale and K. Church, "What is wrong with adding one?" in *Corpus-based research into language*, N. Oostdijk and P. de Haan, Eds. Amsterdam: Rodopi, 1994, pp. 189–98.
- [7] H. E. Robbins, "Estimating the total probability of the unobserved outcomes of an experiment," *Ann. of Math. Stat.*, vol. 39, no. 1, pp. 256–7, 1968.
- [8] B. H. Juang and S. H. Lo, "On the bias of the Turing-Good estimate of probabilities," *IEEE Trans. Signal Processing*, vol. 42, no. 2, pp. 496–8, 1994.
- [9] W. W. Esty, "The efficiency of Good's nonparametric coverage estimator," *Ann. Statist.*, vol. 14, no. 3, pp. 1257–60, 1986.
- [10] C. X. Mao and B. G. Lindsay, "A Poisson model for the coverage problem with a genomic application," *Biometrika*, vol. 89, no. 3, pp. 669–81, 2002.
- [11] D. McAllester and R. E. Schapire, "On the convergence rate of Good-Turing estimators," in *Proc. 13th Annu. Conference on Comput. Learning Theory*. Morgan Kaufmann, San Francisco, 2000, pp. 1–6.
- [12] E. Drukh and Y. Mansour, "Concentration bounds for unigram language models," *J. Mach. Learn. Res.*, vol. 6, pp. 1231–1264, 2005.
- [13] D. McAllester and L. Ortiz, "Concentration inequalities for the missing mass and for histogram rule error," *J. Mach. Learn. Res.*, vol. 4, pp. 895–911, 2003.
- [14] O. Kallenberg, *Foundations of Modern Probability*, 2nd ed. New York: Springer-Verlag, 2002.
- [15] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed. New York: Springer-Verlag, 1998.
- [16] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: John Wiley & Sons, 1991.
- [17] P. Billingsley, *Probability and Measure*, 3rd ed. New York: John Wiley & Sons, 1995.